

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (currently amended) A method for selecting a value in a distributed computing system using a fault tolerant consensus algorithm, the method comprising:

receiving at a computing device from a first client a first message comprising a first proposed value and a first client identifier corresponding to the first client;

provisionally voting at the computing device for the first proposed value;

transmitting from the computing device a first indication of the provisional voting for the first proposed value to one or more devices; ~~and~~

transmitting from the computing device a first result of the provisional voting for the first proposed value to the first client,

wherein the voting for the first proposed value, the transmitting the first indication of the voting for the first proposed value, and the transmitting the first result are not performed if a second message had previously been is received at the computing device from a second client, the second message comprising a second proposed value and a second client identifier corresponding to a second client, the second client identifier being more dominant than the first client identifier, and the second proposed value having been previously provisionally voted for; ~~[[,]] and~~

receiving at the computing device from a third client a third message comprising a third proposed value and a third client identifier corresponding to the third client, the computing device provisionally voting for the third proposed value and transmitting a third indication of the provisional voting for the third proposed value if the third client identifier is dominant as compared to the first client identifier and the third client identifier.

~~wherein the second client identifier corresponding to the second client is determined to be more dominant than the first client identifier by evaluating the second client identifier relative to the first client identifier.~~

2. (Original) The method of claim 1, wherein the first proposed value comprises a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step.

3. (Original) The method of claim 1, wherein the voting for the first proposed value comprises changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier.

4. (Original) The method of claim 1, wherein the first proposed value comprises a first function identified by a first function identifier, and wherein the voting for the first proposed value comprises executing the first function in the first system step unless the first function identifier is equivalent to a second function identifier that identifies a second function, wherein the second function was executed in a second system step that preceded the first system step.

5. (Original) The method of claim 1, wherein the first proposed value comprises a first idempotent function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

6. (Cancelled)

7. (Cancelled)

8. (Original) The method of claim 1 further comprising: transmitting one or more polling messages to initiate a fault tolerant consensus algorithm; receiving one or more vote indication messages in response to the one or more polling messages; and selecting, as a third proposed value, any value if the one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values, or selecting, as the third proposed value, a common possibly selected proposed value if all possibly selected proposed values indicated by the one or more vote indication messages are equivalent to the common possibly

selected proposed value, wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus algorithm.

9. (Currently Amended) A computer-readable storage medium having computer-executable instructions for selecting a value in a distributed computing system using a fault tolerant consensus algorithm, the computer-executable instructions performing steps comprising:

receiving at a computing device from a first client a first message comprising a first proposed value and a first client identifier corresponding to the first client;

provisionally voting at the computing device for the first proposed value;

transmitting from the computing device a first indication of the provisionally voting for the first proposed value to one or more devices; ~~and~~

transmitting from the computing device a first result of the provisional voting for the first proposed value to the first client,

wherein the voting for the first proposed value, the transmitting the first indication of the voting for the first proposed value, and the transmitting the first result are not performed if a second message had previously been ~~is~~ received at the computing device from a second client, the second message comprising a second proposed value and a second client identifier corresponding to the second client, the second client identifier being more dominant than the first client identifier and the second proposed value having been previously voted for, and

receiving at the computing device from a third client a third message comprising a third proposed value and a third client identifier corresponding to the third client, the computing device provisionally voting for the third proposed value and transmitting a third indication of the provisional voting for the third proposed value if the third client identifier is dominant as compared to the first client identifier and the third client identifier

~~wherein the second client identifier corresponding to the second client is determined to be more dominant than the first client identifier by evaluating the second client identifier relative to the first client identifier.~~

10. (Previously Presented) The computer-readable storage medium of claim 9, wherein the first proposed value comprises a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step.

11. (Previously Presented) The computer-readable storage medium of claim 9, wherein the voting for the first proposed value comprises changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier.

12. (Previously Presented) The computer-readable storage medium of claim 11, wherein the second proposed value comprises a second proposed function, and wherein the changing the previous vote comprises undoing a previous execution of the second proposed function.

13. (Previously Presented) The computer-readable storage medium of claim 11, wherein the second proposed value comprises a second proposed function, and wherein the changing the previous vote comprises allowing a previous provisional execution of the second proposed function to expire.

14. (Previously Presented) The computer-readable storage medium of claim 9, wherein the first proposed value comprises a first function identified by a first function identifier, and wherein the voting for the first proposed value comprises executing the first function in the first system step unless the first function identifier is equivalent to a second function identifier that identifies a second function, wherein the second function was executed in a second system step that preceded the first system step.

15. (Previously Presented) The computer-readable storage medium of claim 9, wherein the first proposed value comprises a first idempotent function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

16. (Cancelled)

17. (Cancelled)

18. (Previously Presented) The computer-readable storage medium of claim 9 having further computer-executable instructions for performing steps comprising: transmitting from the computing device one or more polling messages to initiate a fault tolerant consensus method; receiving at the computing device one or more vote indication messages in response to the one or more polling messages; and selecting, as a third proposed value, any value if the one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values, or selecting, as the third proposed value, a common possibly selected proposed value if all possibly selected proposed values indicated by the one or more vote indication messages are equivalent to the common possibly selected proposed value, wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus method.

19. (Currently Amended) A computing device adapted to select a value in ~~operating as~~ ~~part of~~ a distributed computing system using a fault tolerant consensus algorithm, the computing device comprising:

a processing unit performing steps comprising:

comparing at a computer device a first client identifier to a second client identifier if a second proposed value, proposed in a message receives from a second client and comprising the second client identifier and the second proposed value, was previously voted for in a first system step; and

provisionally voting for a first proposed value in the first system step if the first client identifier is more dominant than the second client identifier and the second proposed value was previously voted for;

comparing a third client identifier to a first client identifier and a second client identifier if a third proposed value, proposed in a message received from a third client and comprising the third client identifier and the third proposed value is received; and

a network interface performing steps comprising:

receiving at the computing device from the first client a first message comprising the first proposed value and a first client identifier corresponding to the first client;

transmitting from the computing device a first indication of the voting for the first proposed value to one or more devices also operating as part of the distributed computing system; ~~and~~

transmitting from the computing device a first result of the voting for the first proposed value to the first client;

receiving at the computing device from the third client a third message comprising the third proposed value and a third client identifier corresponding to the third client;

transmitting from the computing device a third indication of the voting for the third proposed value to one or more devices also operating as part of the distributed computing system; and

transmitting from the computing device a third result of the voting for the first proposed value to the third client,

wherein the voting for the first proposed value, the transmitting the first indication of the voting for the first proposed value, and the transmitting the first result are not performed if the second client identifier corresponding to the second client is more dominant than the first client identifier and the second proposed value has been previously voted for, and

wherein the voting for the third proposed value, the transmitting the third indication of the voting for the first proposed value, and the transmitting the third result are performed if the third client identifier corresponding to the third client is more dominant than the first client identifier and the second client identifier

~~wherein the second client identifier corresponding to the second client is determined to be more dominant than the first client identifier by evaluating the second client identifier~~

relative to the first client identifier.

20. (Original) The computing device of claim 19, wherein the first proposed value comprises a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step.

21. (Original) The computing device of claim 19, wherein the voting for the first proposed value comprises changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier.

22. (Original) The computing device of claim 21, wherein the second proposed value comprises a second proposed function, and wherein the changing the previous vote comprises undoing a previous execution of the second proposed function.

23. (Original) The computing device of claim 21, wherein the second proposed value comprises a second proposed function, and wherein the changing the previous vote comprises allowing a previous provisional execution of the second proposed function to expire.

24. (Original) The computing device of claim 19, wherein the first proposed value comprises a first function identified by a first function identifier, and wherein the voting for the first proposed value comprises executing the first function in the first system step unless the first function identifier is equivalent to a second function identifier that identifies a second function, wherein the second function was executed in a second system step that preceded the first system step.

25. (Original) The computing device of claim 19, wherein the first proposed value comprises a first idempotent function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

26. (Previously Presented) The computing device of claim 19, wherein the processing unit performs further steps comprising: ignoring additional proposed values from the first client; and participating in a fault tolerant consensus method; and wherein the network interface performs further steps comprising: receiving a message, wherein the message is part of the fault tolerant consensus method.

27. (Previously Presented) The computing device of claim 26, wherein the participating in the fault tolerant consensus method comprises transmitting a possibly selected proposed value if a proposed value was previously voted for, wherein the possibly selected proposed value was: previously voted for and was proposed by a client having a most dominant client identifier among all clients who proposed values to the computing device for a current system step.

28. (Previously Presented) The computing device of claim 19, wherein the processing unit performs further steps comprising: selecting, as a third proposed value, any value if one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values, or selecting, as the third proposed value, a common possibly selected proposed value if all possibly selected proposed values indicated by the one or more vote indication messages are equivalent to the common possibly selected proposed value, wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus method; and wherein the network interface performs further steps comprising: transmitting one or more polling messages to initiate the fault tolerant consensus algorithm; and receiving the one or more vote indication messages in response to the one or more polling messages.

29. (Original) The computing device of claim 19, wherein the operating as part of the distributed computing system comprises operating as a client of the distributed computing

system.

30. (Original) The computing device of claim 19, wherein the distributed computing system is comprised of devices that are also clients of the distributed computing system.

31. (Currently Amended) A conflict tolerant message delay reducing consensus method for use in a computing environment comprising at least one dedicated client device and a distributed computing system implemented by one or more devices, ~~wherein the computing system implements the conflict tolerant message delay reducing consensus method;~~ the conflict tolerant message delay reducing consensus method comprising:

transmitting one or more proposed values from one or more clients, each of the one or more proposed values being transmitted in a message comprising one of the one or more proposed values and a client identifier corresponding to one of the one or more clients;

voting, at one or more of the one or more devices implementing the distributed computing system, for a proposed value from among the one or more proposed values, wherein the proposed value was proposed by a client having a most dominant client identifier from among the one or more clients proposing values;

transmitting to one or more of the one or more devices implementing the distributed computing system an indication of the vote for the proposed value; and

transmitting, to the client having the highest client identifier, a result of the vote for the proposed value

wherein the voting for the first proposed value, the transmitting the indication of the vote for the proposed value, and the transmitting the result are not performed if the proposed value was proposed by a client not having a most dominant client identifier from among the one or more clients proposing values, ~~and~~

~~wherein the second client identifier corresponding to the second client is determined to be more dominant than the first client identifier by evaluating the second client identifier relative to the first client identifier.~~

32. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31, wherein the one or more devices implementing the distributed

computing system also act as clients of the distributed computing system.

33. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31, wherein the dedicated client device is identified by a least dominant client identifier.

34. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31 further comprising: determining that the distributed computing system has selected the proposed value when each of the one or more devices implementing the distributed computing system has voted for the proposed value.

35. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31, wherein the proposed value comprises a function, and wherein the voting for the proposed value comprises provisionally executing the function in a system step.

36. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31, wherein the voting for the proposed value comprises changing a previous vote if the previous vote was for a previously proposed value, proposed by a previous client having a client identifier that is less dominant than the client proposing the proposed value.

37. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 31 further comprising: ending the conflict tolerant message delay reducing consensus method and commencing a fault tolerant consensus method if a failure is detected.

38. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 37, wherein the failure is detected by a monitoring device, and wherein the computing environment further comprises the monitoring device.

39. (Previously Presented) The conflict tolerant message delay reducing consensus method of claim 37, wherein the commencing the fault tolerant consensus method comprises

DOCKET NO.: MSFT-5033/305892.01
Application No.: 10/750,601
Office Action Dated: November 25, 2009

PATENT

identifying a possibly selected proposed value, wherein the possibly selected proposed value is any value if at least one of the one or more devices implementing the distributed computing system did not previously vote or, if at least one of the one or more devices implementing the distributed computing system previously voted, then the possibly selected proposed value was previously voted for by at least one of the one or more devices implementing the distributed computing system and was proposed by the client having the most dominant client identifier from among the one or more clients proposing values, for a current system step, to the at least one of the one or more devices implementing the distributed computing system.